

INTERNET-DRAFT

Expires: December 2003

Hesham Soliman, Flarion  
Karim ElMalki, Ericsson  
Claude Castelluccia, INRIA  
June, 2003

**Flow movement in Mobile IPv6**  
<[draft-soliman-mobileip-flow-move-03.txt](#)>

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This document is an individual submission to the IETF. Comments should be directed to the authors.

Abstract

The aim of this draft is to introduce a new extension to MIPv6 to allow hosts to direct inbound flows individually to certain preferred interfaces. This extension to MIPv6 allows multihomed hosts to take full advantage of the diverse access technologies that they may be connected to and direct their traffic according to internal policies specified by the users or applications.



## 1. Introduction

The current MIPv6 specification [[MIPv6](#)] allows a MN to manage its CoA by sending BUs to its HA and other CNs when applicable. The semantics of the BUs in MIPv6 are limited to host movement. I.e. The current MIPv6 specification does not allow a MN to split its inbound connections to different addresses. In this draft, the splitting of inbound traffic to be received on different addresses is referred to as "Per-flow movement".

In the context of this proposal, a flow can be defined as one or more connections that are identified by a flow identifier. A single connection is typically identified by the source and destination IP addresses, transport protocol number and the source and destination port numbers. Alternatively a flow can be identified in a simpler manner using the flow label field in the IPv6 header [[IPv6](#)].

Flow movement can be a useful feature in cases where the MN is connected to different access technologies with different characteristics. When using the flow movement options below, a MN would be able to "move" one flow to an interface while maintaining the reception of other flows on another interface. Requesting the flow movement can be decided based on local policies within the MN and based on the link characteristics and the types of applications running at the time.

It should be noted that the flow movement option can be associated with any BU, whether it is sent to a CN, HA or MAP [[HMIPv6](#)]. A similar mechanism for Mobile IPv4 is described in [[FNS01](#)].

## 2. Flow movement option

The Flow movement options are included within the BU and BA messages. These options contain information that allows the receiver of a BU to identify a traffic flow and route it to a given address. Multiple options may exist within a BU. These options may contain the same destination IPv6 address or different addresses. Only one destination address is allowed in each option.

A traffic flow may be identified by using the flow label in IPv6 or by combining the destination address, transport protocol number and port number. Two different types of options are defined in this memo, one identifies a flow based on the addresses/protocol number/port numbers quintuplet, and the other identifies the connection based on the flow label combined with the CN's source address.

A MN can include several options within the BU message. For instance, a MN could move a number of connections to another interface. In the absence of a defined mechanism for flow label usage the MN would

include a number of flow movement options, each identifying one

Soliman, El-Malki, Castelluccia

[Page 2]

connection based on the source/destination addresses, port numbers and the protocol number quintuplet.

It should be noted that per-packet load balancing has negative impacts on TCP congestion avoidance mechanisms as it is desirable to maintain order between packets belonging to the same TCP connection. This behaviour is specified in [TRAFF]. Other negative impacts are also foreseen for other types of real time connections due to the potential variations in RTT between packets. Hence per-packet load balancing is not allowed in this extension. However, the MN can still request per-flow load balancing provided that the entire flow is moved to the new interface.

**2.1 Option format for flow classification based on port numbers**

Figure 1 shows the option format used when using the addresses/protocol number/ port numbers quintuplet to classify a flow. The MN's destination address, to which the flow is being moved, is assumed to be the source address in the IP header. Hence, when using this mechanism, the MN MUST use the appropriate source address in the IP header.

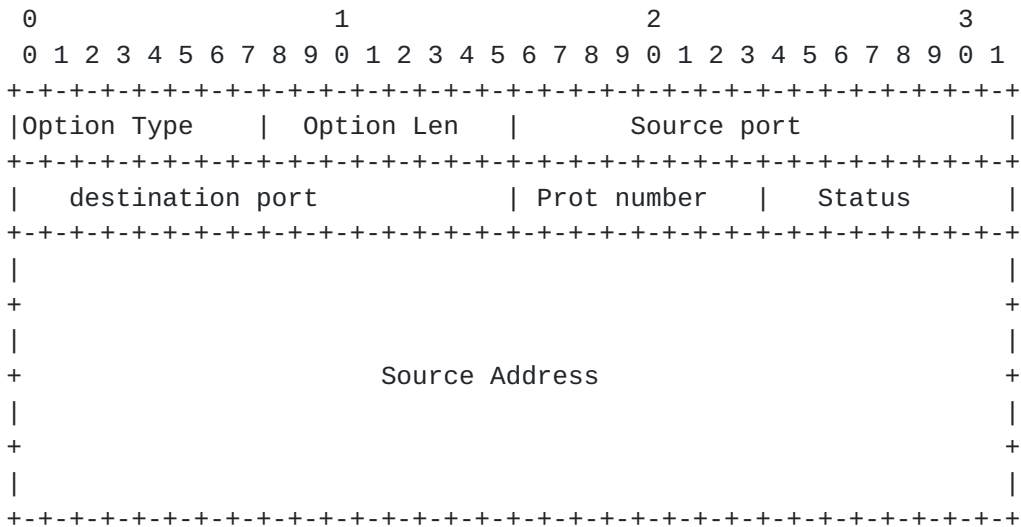


Figure 1. Port numbers based flow selection

Option Type	TBD
Option Len	Length of option
Source port	The port number for the CN
Destination port	The port number for the MN

Prot number

A 16-bit unsigned integer representing

Soliman, El-Malki, Castelluccia

[Page 3]

value of the transport protocol number associated with the port numbers.

Status                    An unsigned 8 bit integer indicating the success or failure for this option. Values lower than 128 are reserved for successful registrations. Failure values are 128 and above. This field is used to indicate the success or failure of the operation when the option is part of the BA. It is also used in the BU to indicate whether the option should be added to, or deleted from, the binding cache. When set to Zero, it indicates addition, and a value Of 0xFF indicates a request for deletion (deregistration).

The following values are reserved for the status field within the flow movement option:

- 0     Indicates a successful registration.
- 128   Flow movement rejected, reason unspecified.
- 129   Flow movement option poorly formed.
- 130   Flow identification by port numbers is not Supported.

Source Address            A 128-bit field representing the source Address of the CN.

The alignment requirement for this option is 8n.

## **2.2 Option format for flow classification based on the Flow label**

Figure 2 shows the option format for flow splitting based on the Flow label and the source address. As mentioned above, the MN's destination address is assumed to be the source address in the IP header, hence the MN MUST select the source address in light of this requirement.





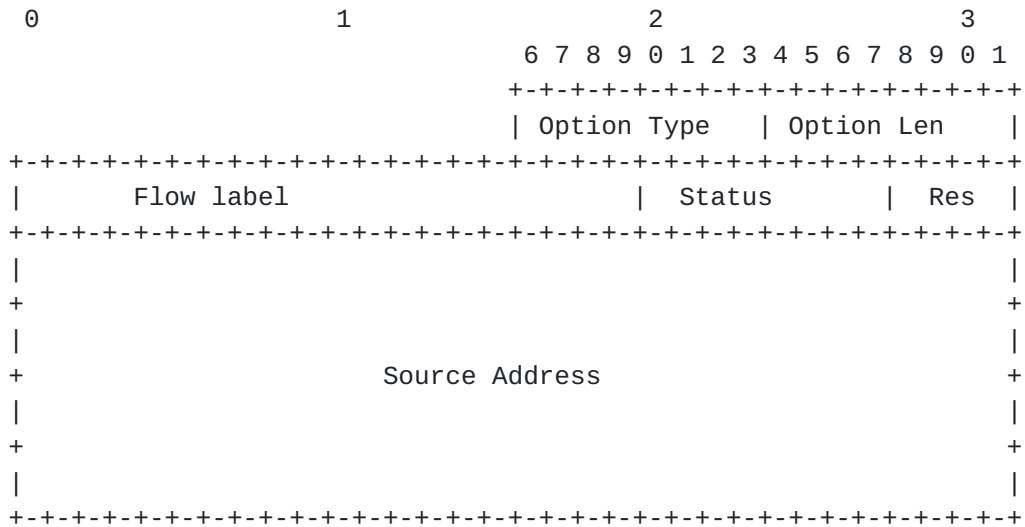


Figure 2. Flow label based flow selection

Option Type	TBD
Option Len	Length of option
Status	An unsigned 8 bit integer indicating the success or failure for this option. Values lower than 128 are reserved for successful registrations. Failure values are 128 and above. This field is only used when the option is part of the BA to indicate the operation's success or failure. It is also used in the BU to indicate whether the option should be added to, or deleted from, the binding cache. When set to Zero, it indicates addition, and a value Of 0xFF indicates a request for deletion (deregistration).

The following values are reserved for the status field within the flow movement option:

- 0 Indicates a successful registration.
- 128 Flow movement rejected, reason unspecified.
- 129 Flow movement option poorly formed.
- 130 Flow identification by flow label is not Supported.

Res	A 4-bit reserved field, MUST be set to Zero by the sender and ignored by the
-----	--

Receiver.

Soliman, El-Malki, Castelluccia

[Page 5]

Source Address            A 128-bit field representing the source Address of the CN.

### **3. Sending rules for the MN**

For this mechanism to be useful, the MN MUST ensure that the appropriate Source address (for the CN) is used in the option. This is clear when sending the BU directly to the CN, as both ends possess the necessary information required to identify the connection.

However, when the BU is sent to an intermediate router, like the HA or MAP, careful selection of the CN's source address is required. The reason for this is that the CN may also be a MN. The remaining part of this section will consider the case where the MN is sending BUS to an intermediate router, like a HA or MAP.

If the CN is not a MN, the source address can be assumed to be the destination address that the MN's applications use to send traffic to the CN. This implies that the source address field in the flow-movement option is the same address that the MN uses as part of the quintuple identifying the connection (i.e. the destination address for the connection, seen by upper layers).

However, if the CN is also a MN, sending BUs, the CN's address is the CoA stored in the MN's binding cache. This is the source address included in the IPv6 header seen by intermediate nodes.

### **4. Deregistering the Flow movement option**

A MN may, at some point in time, decide to deregister the Flow movement option due to connection termination or a change in its IP layer access point. This can be achieved by resending the BU with the status field set to 0xFF.

### **5. Acknowledging the Flow movement option**

The receiver of the Flow movement option MUST acknowledge it in a way that allows the sender to maintain the option's information in its binding update list. If one or more options are accepted, the CN MUST include all the options with the appropriate Status values in the BA.

The acceptance of each flow movement option is independent from the acceptance of the CoA in the BU as well as other options. In other words, the acceptance of the new CoA in a BU does not imply an acceptance of every flow movement option. Hence, the receiver of the BU MUST include all the flow movement options in the BA with an appropriate status value to indicate the acceptance or rejection of each one. This will ensure consistency in the Binding Cache of the

receiver and the BU list of the sender. If the receiver of the flow

movement option does not include it in its BA with an appropriate Status code, the sender should assume that the option was not accepted.

### **5.1 Additional Binding Acknowledgement status values**

A New BA status value will need to be introduced to support the flow movement feature. The new value is shown below:

1 Binding Update accepted, flow movement is not supported.

This implies the rejection of all the Flow movement options. If this code is used, the CN SHOULD NOT include any of the Flow movement options in the reply.

## **6. Notice regarding Intellectual Property Rights**

see <http://www.ietf.org/ietf/IPR/ERICSSON-General>

## **7. Acknowledgements**

A Special acknowledgement goes to Wolfgang Hansmann for his careful reviews and suggestions to improve this draft. Thanks to Conny Larsson for his review of the draft and helpful Comments, and to: Simon Aladdin, Tomas Goransson as well as other members of the DRiVE project for their useful input towards this draft.

## **8. References**

- [FNS01] X.Zhao, C.Castelluccia and M.Baker. "Flexible Network Support for Mobile Hosts", ACM MONET, April 2001.
- [HMIPv6] H. Soliman, C. Castelluccia, K. ElMalki and L. Bellier "Hierarchical MIPv6 mobility management". [draft-ietf-mobileip-hmipv6-05.txt](#)
- [MIPv6] D. Johnson and C. Perkins, "Mobility Support in IPv6", [draft-ietf-mobileip-ipv6-13.txt](#), February 2000.
- [IPv6] S. Deering and B. Hinden, "Internet Protocol version 6 (IPv6) specification". [RFC 2460](#).
- [TRAFF] D. Awduche et al, "Requirements for traffic engineering over MPLS". [RFC 2702](#).

## **9. Authors' addresses**

Hesham Soliman  
Flarion Technologies

E-mail: H.Soliman@Flarion.com

Soliman, El-Malki, Castelluccia

[Page 7]

Karim El Malki  
Ericsson Radio Systems AB  
Access Networks Research  
SE-164 80 Stockholm  
SWEDEN

Phone: +46 8 7195803  
Fax: +46 8 7575720  
E-mail: Karim.El-Malki@era.ericsson.se

Claude Castelluccia  
INRIA /Planete  
ZIRST- 655 Avenue de l'Europe  
38334 Saint Ismier Cedex  
France

